arXiv:1601.02149v1 [q-fin.PR] 9 Jan 2016

Computing semiparametric bounds on the expected payments of insurance instruments via column generation*

Robert Howley†[1], Robert Storer‡[1], Juan Vera§[2] and Luis F. Zuluaga¶[1]

[1] *Department of Industrial and Systems Engineering, Lehigh University*

H.S. Mohler Laboratory, 200 West Packer Avenue, Bethlehem, PA 18015

[2] *Department of Econometrics and Operations Research, Tilburg University*

5000 LE TILBURG The Netherlands

Abstract

It has been recently shown that numerical semiparametric bounds on the expected payoff of financial or actuarial instruments can be computed using semidefinite programming. However, this approach has practical limitations. Here we use column generation, a classical optimization technique, to address these limitations. From column generation, it follows that practical univariate semiparametric bounds can be found by solving a series of linear programs. In addition to moment information, the column generation approach allows the inclusion of extra information about the random variable; for instance, unimodality and continuity, as well as the construction of corresponding worst/best-case distributions in a simple way.

1 Introduction

Many financial and insurance instruments protect against underlying losses for which it is difficult to make exact distributional assumptions. Under these circumstances, it is difficult to provide a good estimate of the loss distribution, which in turn makes it difficult to estimate payments on the corresponding insured loss. Computing *semiparametric bounds* on the expected payments is an

*This work has been carried out thanks to funding from the Casualty Actuarial Society, Actuarial Foundation of Canada, and Society of Actuaries

†howley.robert@gmail.com

‡rhs2@lehigh.edu

§j.c.veralizcano@uvt.nl

¶luis.zuluaga@lehigh.edu, Corresponding author.

approach that has been successfully used to deal with this problem. This involves finding the minimum and maximum expected payments on the insurance instrument, when only partial information (e.g., moments) of the underlying loss distribution is known. For example, consider the work of Cox (1991); Jansen et al. (1986); Villegas et al. (2010). This approach has also been used to address the estimation of bounds on extreme loss probabilities (Cox et al. (2010)), and the prices of insurance instruments, and financial options (Brockett et al. (1996); Lo (1987); Schepper and Heijnen (2007)). These semiparametric bounds are useful when the structure of the product is too complex to develop analytical or simulation based valuation methods, or when it is difficult to make strong distributional assumptions on the underlying risk factors. Furthermore, even when distributional assumptions can be made, and analytical valuation formulas or simulation based prices can be derived, these bounds are useful to check the consistency of such assumptions.

The semiparametric bound approach is also referred as *distributionaly-robust* (see, e.g., Delage and Ye, 2010) or *ambiguity-averse* (see, e.g., Natarajan et al., 2011). Also, it has been shown that this approach partially reflects the manner in which persons naturally make decisions (cf., Natarajan et al., 2011)

In the actuarial science and financial literature, there are two main approaches used to compute semiparametric bounds: analytically, by deriving closed-form formulas for special instances of the problem (see, e.g., Chen et al., 2011; Cox, 1991; Schepper and Heijnen, 2007); and numerically, by using *semidedefinite programming* techniques (cf., Todd, 2001) to solve general instances of the problem (see, e.g., Bertsimas and Popescu, 2002; Boyle and Lin, 1997; Cox et al., 2011, 2010). An alternative numerical approach to solve semiparametric bounds proposed in the *stochastic programming* literature by Birge and Dulá (1991), based on the classical *column generation* (CG) approach for mathematical optimization problems (see, e.g., Dantzig (1963, Chp. 22), Lubbecke and Desrosiers (2005)), has received little attention in the financial and actuarial science literature.

Here, we consider the use of CG to obtain semiparametric bounds in the context of financial and actuarial science applications. In particular, we show that for all practical purposes, univariate semiparametric bounds can be found by solving a sequence of linear programs associated to the CG *master* problem (cf., Section 3). We also show that the CG approach allows the inclusion of extra information about the random variable such as unimodality and continuity, as well as the construction of the corresponding worst/best-case distributions in a simple way. Also, the CG methodology

achieves accurate results at a very small computational cost, it is straightforward to implement, and the core of its implementation remains the same for very general, and practical instances of semiparametric bound problems.

To illustrate the potential of the CG approach, in Section 5.1, semiparametric lower and upper bounds are computed for the loss elimination ratio of a right censored deductible insurance policy, when the underlying risk distribution is assumed to be unimodal, and have known first and second-order moments. Furthermore, in Section 5.2, we illustrate how continuous representations of the worst/best-case distributions associated with the semiparametric bounds can be readily constructed and analyzed.

2 Problem Description

Consider a random variable X with an unknown underlying distribution π, but known support $\mathcal{D} \subseteq \mathbb{R}$ (not necessarily finite), and interval estimates $[\sigma_j^-, \sigma_j^+]$, $j = 1, \ldots, m$ for the expected value of functions $g_j : \mathbb{R} \to \mathbb{R}$ for $j = 1, \ldots, m$ (e.g., typically, $g_j(x) = x^j$). The *upper semiparametric bound* on the expected value of the (target) function $f : \mathbb{R} \to \mathbb{R}$ is defined as:

$$\begin{aligned} B^* := \sup_{\pi} \quad & E_\pi\left(f(X)\right) \\ \text{s.t.} \quad & \sigma_j^- \leq E_\pi\left(g_j(X)\right) \leq \sigma_j^+ \qquad j = 1, \ldots, m \\ & \pi \text{ a probability distribution on } \mathcal{D}, \end{aligned} \tag{1}$$

where $E_\pi(\cdot)$ represents the expected value under the distribution π. That is, the upper semiparametric bound of the function f is calculated by finding the supremum of $E_\pi\left(f(X)\right)$ across all possible probability distributions π, with support on the set $\mathcal{D}$, that satisfy the $2m$ expected value constrains. The parameters σ_j^-, σ_j^+, $j = 1, \ldots, m$ allow for confidence interval estimates for the expected value of $E_\pi(g_j(X))$, that are not typically considered in the analytical solution of special instances of (1) (i.e., typically $\sigma_j^- = \sigma_j^+$ in analytical solutions).

The *lower semiparametric bound* of the function f is formulated as the corresponding minimization problem, that is, by changing the sup to inf in the objective of (1). We will provide details about the solution of the upper semiparametric bound problem (1) that apply in analogous fashion to the corresponding lower semiparametric bound problem. Also, for ease of presentation we will at times

refer to both the upper and lower bound semiparametric problems using (1).

While specific instances of (1) have been solved analytically (see, e.g., Lo, 1987; Cox, 1991; Schepper and Heijnen, 2007; Chen et al., 2011), semidefinite programming (SDP) is currently the main approach used in the related literature to numerically solve the general problem being considered here (c.f., Boyle and Lin, 1997; Bertsimas and Popescu, 2002; Popescu, 2005) whenever the functions $f(\cdot)$ and $g_j(\cdot)$ are piecewise polynomials. However, the SDP approach has important limitations in terms of the capacity of practitioners to use it. First, there are no commercially available SDP solvers. Second, the formulation of the SDP that needs to be solved for a given problem is not "simple" (Cox et al. (2010)) and must be re-derived for different support sets $\mathcal{D}$ of the distribution of X (see, e.g., Bertsimas and Popescu, 2002, Proposition 1).

Birge and Dulá (1991) proposed an alternative numerical method to solve the semiparametric bound problem (1) by using a CG approach (see, e.g., Dantzig (1963, Chp. 22), Lubbecke and Desrosiers (2005)) that has received little attention in the financial and actuarial science literature. Here, we show that the CG solution approach addresses the limitations of the SDP solution approach discussed above. Additional advantages of the CG solution approach in contrast to SDP techniques will be discussed at the end of Section 3.

It is worth to mention that although in the next section we present the proposed algorithm in pseudo-algorithmic form (e.g., see Algorithm 1), our implementation of the algorithm is available upon request to the authors.

3 Solution via column generation

In this section we present the CG solution approach proposed by Birge and Dulá (1991, Sec. 3) to solve the semiparametric bound problem (1). For the sake of simplifying the exposition throughout we will assume that (1) has a feasible solution, and that the functions $f(\cdot), g_j(\cdot),\ j = 1, \ldots, m$ are Borel measurable in $\mathcal{D} \subseteq \mathbb{R}$ (cf., Zuluaga and Peña, 2005). Now let $J \subseteq \mathcal{D}$ be a set of given *atoms*, and construct the following linear program (LP) related to (1) by associating a probability

decision-variable p_x for every $x \in J$:

$$
\begin{array}{lll}
M_J^* := & \max_{p_x} & \displaystyle\sum_{x\in J} p_x f(x) \\
& \text{s.t.} & \sigma_j^- \le \displaystyle\sum_{x\in J} p_x g_j(x) \le \sigma_j^+ \qquad j = 1,\ldots,m, \\
& & \displaystyle\sum_{x\in J} p_x = 1, \\
& & p_x \ge 0 \qquad\qquad\qquad\qquad \text{for all } x \in J
\end{array}
\tag{2}
$$

Furthermore, we assume that the set $J \subseteq \mathcal{D}$ is *feasible*; that is, the corresponding LP (2) is feasible. The existence of such $J \subseteq \mathcal{D}$ follows from the classical result by Kemperman (1968, Theorem 1), and can be found by solving algorithmically a *Phase I* version (cf., Bertsimas and Tsitsiklis, 1997) of the CG Algorithm 1. Following CG terminology, given a set $J \subseteq \mathcal{D}$ we will refer to (2) as the *master* problem.

Notice that any feasible solution of problem (2) will be a feasible (atomic distribution) for problem (1). Also, the objectives of the two problems are the expected value of the function $f(x)$ over the corresponding decision variable distribution. Thus, M_J^* is a lower bound for the optimal value of the upper semiparametric bound problem (1). Furthermore, it is possible to iteratively improve this lower bound by updating the set $J \subseteq \mathcal{D}$ using the optimal *dual values* (cf., Bertsimas and Tsitsiklis, 1997) of the constrains after the solution of the master problem (2). Namely, let ρ_j^-, ρ_j^+, $j = 1,\ldots,m$ and τ be the dual variables of the upper/lower moment (i.e., first set of constraints in eq. (2)) and total probability (i.e, $\sum_{x\in J} p_x = 1$) constrains respectively. Given a feasible set $J \subseteq \mathcal{D}$, the dual variables can be used to select a new point $x \in \mathcal{D}$, to add to $J \subseteq \mathcal{D}$, that will make the corresponding LP (2) a tighter approximation of (1). In particular, given ρ_j^-, ρ_j^+, $j = 1,\ldots,m$ and τ, consider the following *subproblem* to find x.

$$
S_{\rho,\tau}^* := \max_{x\in\mathcal{D}} \quad f(x) - \tau - \sum_{j=1}^{m} (\rho_j^+ + \rho_j^-) g_j(x) \tag{3}
$$

The objective value of (3) represents the *reduced cost* of adding the new point x to J; that is, the marginal amount by which the objective in (2) can be improved with the addition of x in the master problem (2). Using the master problem (2) and subproblem (3) admits an iterative algorithm that (under suitable conditions) converges to the optimal value of (1). More specifically, at each iteration,

the master problem (2) is solved and its corresponding dual variables are used in the subproblem (3) to select a new point x to be added to the set $J \subseteq \mathcal{D}$. This is called a CG algorithm since at each iteration, a new variable p_x, corresponding to the new given point x, is added to the master problem (2).

If the functions $f(\cdot)$, $g_j(\cdot), j = 1, \ldots, m$ are continuous, and the support $\mathcal{D}$ of the underlying risk distribution is known to be compact, then the asymptotic convergence of the column generation algorithm follows from Dantzig (1963, Thm. 5, Chp. 24). However, for the numerical solution of the practical instances of (1) considered here, it suffices to have a "stopping criteria" for the CG algorithm.

Theorem 1. *Let $J \subseteq \mathcal{D}$ be given, and B^*, M_J^*, $S_{\rho,\tau}^*$ be the optimal objective values of* (1), (2), *and* (3) *respectively. Then* $0 \leq B^* - M_J^* \leq S_{\rho,\tau}^*$.

Theorem 1 follows from Dantzig (1963, Thm. 3, Chp. 24), and states that the LP approximation from below in (2) will be within ϵ of the optimal objective of (1) if the objective of subproblem (3) is less than ϵ. It is worth mentioning that under additional assumptions about the feasible set of (1), one has that in the long-run $S_{\rho,\tau}^* \to 0$; that is, the CG algorithm will converge to the optimal solution of the semiparametric bound problem (1). This follows as a consequence of Dantzig (1963, Thm. 5, Chp. 24). In practice, Theorem 1 provides a stopping criteria for the implementation of the CG algorithm under only the assumption of the original problem (1) being feasible. Specifically, the CG Algorithm 1 can be used to find the optimal upper bound B^* up to ϵ-accuracy. As mentioned before, a *Phase I* version (cf., Bertsimas and Tsitsiklis, 1997) of the CG Algorithm 1, can be used to construct an initial feasible set $J_0 \in \mathcal{D}$.

Algorithm 1 Semiparametric bounds via column generation

```
procedure GC(feasible J_0, ε > 0)
    J ← J_0, S*_{ρ,τ} = ∞
    while S*_{ρ,τ} > ε do
        compute M*_J, p* := {p*_x}_{x∈J}, the optimal objective and solution of master problem (2)
        compute S*_{ρ,τ}, x* , the optimal objective and solution of subproblem (3)
        J ← J ∪ {x*}
    end while
    return J* = J, p*, and M*_J ≈ B* (where ≈ stands for M*_J approximates B* )
end procedure
```

Note that (in principle) problem (1) has an infinite number of columns (i.e., variables) and a finite

number of constrains; that is, the semiparametric bound problem (1) is a *semi-infinite* program. Thus, the approach outlined above is an application to a semi-infinite program of column generation techniques initially introduced by Dantzig and Wolfe (1961) for LPs, generalized linear programs, and convex programs (cf. Dantzig (1963, Chp. 22–24)). For a survey of column generation methods see Lubbecke and Desrosiers (2005).

3.1 Solving the subproblem

As observed by Birge and Dulá (1991), the main difficulty in using the CG approach to solve the semiparametric bound problem (1) is that the subproblem (3) is in general a *non-convex* optimization problem (cf., Nocedal and Wright, 2006). However, in the practically relevant instances of the problem considered here, the following assumption holds.

Assumption 1. *The functions $f(\cdot)$, and $g_j(\cdot)$, $j = 1, \ldots, m$ in (1) are piecewise polynomials of degree less than five (5).*

More specifically, typically no higher than fifth-order moment information on the risk will be assumed to be known (e.g., $g_j(X) = X^j$ for $j = 1, \ldots, m$ and $m \leq 5$). Also, the function $f(\cdot)$ typically defines: the piecewise linear cost or payoff of an insurance instrument (e.g., $f(x) = \max\{0, x - d\}$); a ruin event using a (piecewise constant) indicator function (e.g., $f(x) = \mathbb{I}_{[0,r]}(x)$) ; or a lower than fifth-order moment of the risk or insurance policy cost/payoff (see, e.g. Brockett et al., 1996; Cox, 1991; Lo, 1987; Schepper and Heijnen, 2007). In such cases, the objective of (3) is an univariate fifth-degree (or lower) piecewise polynomial. Thus, subproblem (3) can be solved "exactly" by finding the roots of fourth degree polynomials (using *first-order* optimality conditions (see, e.g., Nocedal and Wright, 2006)). As a result, we have the following remark.

Remark 2. *Under Assumption 1 and using the CG Algorithm 1, the solution to problem (1) can be found by solving a sequence of LPs (2) where the column updates (3) can be found with simple arithmetic operations.*

Moreover, thanks to current numerical algorithms for finding roots of univariate polynomials, it is not difficult to solve subproblem (3) numerically to a high precision, even when the polynomials involved in the problem have degree higher than 5. In turn, this means that Algorithm 1 would perform well for instances of the problem with high degree polynomials.

Generating semiparametric bounds using the CG approach outlined above has several key advantages over the semidefinite programming (SDP) solution approach introduced by Bertsimas and Popescu (2002); Popescu (2005). First, only a linear programming solver for (2) and the ability to find the roots of polynomials with degree no more than four for (3) is required in most practically relevant situations. This means that the methodology can use any commercial LP solver allowing for rapid and numerically stable solutions. Second, the problem does not need to be reformulated for changes in the support $\mathcal{D}$ of the underlying risk distribution π of X. Accounting for alternate support requires only limiting the search space in the subproblem (3). Finally, problem (2) is explicitly defined in terms of the distribution used to generate the bound value. So, for any bound computed, the *worst-case* (resp. *best-case* for the lower bound) distribution that yielded that bound can also be analyzed; with the SDP approach no such insight into the resulting distribution is, to the best of our knowledge, readily possible. The ability to analyze the resulting distribution would be of particular use to practitioners in the insurance and risk management industry and will be further discussed in Section 5.2. Third, the CG approach works analogously for both the upper and lower semiparametric bound problems. In contrast, the SDP approach commonly results in SDP formulations of the problem that are more involved for the lower than for the upper semiparametric bound. Finally, as shown in Section 4, the CG approach allows the addition of information about the class of distribution to which the underlying risk belongs (e.g., continuous, unimodal) without changing the core of the solution algorithm.

4 Additional Distribution Information

As mentioned earlier, in practical instances of the semiparametric bound problem (1) the functions $g_j(\cdot),\ j = 1, \ldots, m$ are typically set to assume the knowledge of moments of the underlying loss distribution; for example, by setting $g_j(X) = X^j,\ j = 1, \ldots, m$ in (1). The general semiparametric bound problem (1) can be extended to include additional distributional information other than moments (see, e.g., Popescu, 2005; Schepper and Heijnen, 2007). This is important as the resulting bounds will be tighter and the corresponding worst/best-case distribution will have characteristics consistent with the practitioner's application-specific knowledge about continuity, unimodality, and heavy tails in financial loss contexts. In this section it is shown that the CG solution approach

outlined in Section 3 for semiparametric bound problems can be extended to constrain the underlying distribution to be unimodal and continuous.

4.1 Mixture Transformation

Note that a point $x \in J^*$ obtained after running the CG Algorithm 1 can be interpreted as the mean of Dirac delta distributions δ_x parametrized by (centered at) x. In turn, the resulting optimal distribution π^* of the random variable X in (2) is a mixture of Dirac delta distributions; that is, $\pi^* \sim \sum_{x \in J^*} p_x \delta_x$. As we show below, the CG Algorithm 1 can be used to obtain optimal worst/best-case distributions associated with the semiparametric bound problem (1) when besides the expected value constrains, information is known about the class of distribution to which π belongs; for example, unimodal, smooth, asymmetric, etc. Basically this is done by replacing $\delta_x \to H_x$ in the mixture composition of π, where H_x is an appropriately chosen distribution parametrized by x.

More specifically, assume that besides the moment information used in the definition of the semiparametric bound problem (1), it is known that the distribution π is a mixture of known probability distributions H_x, parametrized by a single parameter $x \in \mathbb{R}$. For example, x could be the mean of the distribution H_x, or H_x could be a uniform distribution between 0 and x. Note that for any $g : \mathbb{R} \to \mathbb{R}$, it follows from the mixture composition of the distribution π that:

$$\begin{aligned} E_\pi(g) &= \int_0^\infty g(u) E_{\pi(X)}(H_X(u)) du \\ &= E_{\pi(X)} \left(\int_0^\infty g(u) H_X(u) du \right) \\ &= E_{\pi(X)} \left(E_{H_X(U)}(g(U)) \right). \end{aligned} \tag{4}$$

This means that with the additional distribution mixture constrain, the associated semiparametric bound problem can be solved with the CG Algorithm 1 after replacing

$$\begin{aligned} f(x) &\to \int_0^\infty f(u) H_x(u) du = E_{H_x}(f), \\ g_j(x) &\to \int_0^\infty g_j(u) H_x(u) du = E_{H_x}(g_j), \end{aligned} \tag{5}$$

for $j = 1, \ldots, m$ in (2), and (3).

Note that in many instances, the expectations in (5) can be computed in closed-form as a function

of the mixture distribution parameter x. Moreover, the expectations in (5) are commonly piecewise polynomials in x (e.g., if H_x is a uniform distribution between 0 and x), or can be written as polynomials after an appropriate change in variable (e.g., if H_x is a lognormal distribution with chosen volatility parameter $\sigma \in \mathbb{R}^+$, and mean $e^{x+\frac{1}{2}\sigma^2}$). In such cases, after applying the transformation (5) the objective of subproblem (3) will be a piecewise polynomial on x. As discussed in Section 3.1, the subproblem can then be solved "exactly", and Remark 2 will still hold as long as Assumption 1 is valid after the transformation (5). This is illustrated in Example 1 below. Also, as we show with numerical experiments in Section 5, this is the case in most practical applications.

Example 1. *Consider a simple insurance policy with no deductilble on a loss X for which the non-central moments up to m-order are assumed to be known. Specifically, let $f(x) = \max\{0, x\}$, and $g_j(x) = x^j$, $j = 1, \ldots, m$. Also, assume that the distribution of the loss X is known to be a mixture of uniform distributions H_x of the form $H_x \sim \text{Uniform}(0, x)$ in (4). That is, $H_x(u) = \frac{1}{x}\mathbb{I}_{[0,x]}(u)$. From (5), it follows that $E_{H_x}(f) = \frac{x}{2}$, $E_{H_x}(g_j) = \frac{1}{j+1}x^j$ for $j = 1, \ldots, m$, and Remark 2 will hold for any $m \leq 4$.*

In other cases; that is, when Assumption 1 does not hold after the transformation (5), one can sharply approximate the expectations in (5) using up to fifth-degree piecewise polynomials on x to take advantage of Remark 2. Alternatively, given that the subproblem (3) is an univariate optimization problem, global optimization solvers such as `BARON` (cf., Tawarmalani and Sahinidis (2005)) can be used to effectively solve it.

In Section 4.2 we will discuss how the mixture transformation (5) can be used to substantially strengthen semiparametric bounds by using reasonable assumptions about the underlying risk distribution regarding unimodality, and or continuity by using a mixture of appropriate distributions. Moreover, in Section 5.2, we use this transformation to construct reasonable worst/best-case distributions associated to a given semiparametric bound problem.

4.2 Unimodality

In many instances of the semiparametric bound problem, it might be reasonable to assume that the unknown distribution π of X in (1) is unimodal with known mode M . This is particularly the case when the underlying random variable represents a financial asset or a portfolio of financial assets

which are typically modelled by a lognormal distribution when using parametric techniques (see, e.g. Schepper and Heijnen, 2007). In this section, we discuss how the unimodality information can be used in a straightforward fashion within the CG algorithm solution approach to obtain tighter semiparametric bounds. Before discussing this below, it is worth mentioning that there are problems in the context to Actuarial Science where it is not appropriate to assume unimodality. For example, as shown in Lee and Lin (2010), this is the case when the underlying random variable is associated with Property/Casualty Losses which often exhibit a multimodal behaviour due to the combination sources compounding the loss (e.g., fire, wind, storms, hurricanes).

It has been shown by Popescu (2005) that semiparametric bound problems with the additional constrain of the underlying distribution being unimodal can also be reformulated as a SDP by calling upon the classical probability result by Khintchine (1938) regarding the representation of unimodal distributions. Specifically, Khintchine (1938) proved that any unimodal distribution can be represented by a mixture of uniform distributions each of which have M as an endpoint (either the right or left endpoint). This same result can be embedded in the framework of Section 3 by leveraging the variable transformation of Section 4.1.

Recall that the CG algorithm can be defined in terms of mixing distributions H_x, where x represents a parameter of the distribution. In particular, for given (mode) $M \in \mathbb{R}$, let

$$H_x \sim \text{Uniform}(\min\{x, M\}, \max\{x, M\}). \tag{6}$$

Using H_x above in (5) to transform the semiparametric bound problem (1) will lead to a bound over distributions that are unimodal with mode M.

The simple transformation (5) using the mixture of uniforms (6) allows the CG approach to leverage the results of Khintchine (1938) while avoiding a complex reformulation as in the case of the SDP methodology of Popescu (2005). Enforcing unimodality is a straightforward special case that highlights the flexibility of the methodology discussed in Section 3.

4.3 Smoothness and Unimodality

The base method of Section 3 computes the desired semiparametric bounds, and provides a discrete (atomic) worst/best-case distribution (x, p_x) for all $x \in J$ associated with the bound. In practice

it is more desirable and intuitive to work with a continuous probability density function. If one is considering a problem measuring financial loss, then having discrete loss values may not provide the insight that a continuous probability density function would, given that a discrete collection of outcomes is highly unrealistic. Using the uniform mixture defined in (6) is guaranteed to yield a unimodal distribution in the computation of the semiparametric bounds (1). However, the resulting density will contain multiple discontinuities including at the mode itself. Furthermore, the density will only be nonzero over the interval $[\min\{x : x \in J^*\}, \max\{x : x \in J^*\}]$; that is, it has finite support. It would be desirable to obtain worst/best-case distributions associated with the semiparametric bounds that are *smooth*; that is, both continuous and differentiable.

Below we show that by appropriately choosing the distribution H_x (and its parameters) in the mixture, it is possible to obtain worst/best-case distributions that are both smooth and unimodal, and that closely replicate the corresponding upper (best) and lower (worst) semiparametric bounds. This can be readily done using the CG approach by reformulating the semiparametric bound problem (1) using the transformation (5), and choosing

$$H_x \sim \text{lognormal}(\mu_x, \sigma_x), \tag{7}$$

where μ_x, σ_x are given in terms of x by the equations

$$\begin{aligned} e^{\mu_x + \frac{1}{2}\sigma_x^2} &= x, \\ (e^{\sigma_x^2} - 1)e^{2\mu_x + \sigma_x^2} &= \alpha^2. \end{aligned} \tag{8}$$

for a given $\alpha \in \mathbb{R}^+$. That is, the lognormal distribution H_x is set to have a mean of x, and variance α^2. Note that besides the mean parameter x, which will be used to construct the mixture using the CG Algorithm 1, one needs to set a second parameter α in (8) to properly define the lognormal distribution H_x in (7).

The lognormal mixture approach (i.e., (7), and (5)) can be used to obtain solutions to the semiparametric bound problem (1) where the underlying worst/best-case distribution is both unimodal, and smooth, and replicates as close as possible the semiparametic bound obtained when the distribution is assumed to be unimodal (and not necessarily smooth). This is in part thanks to the additional

degree of freedom given by the choice of the parameter α in (8). To see this, let us refer to

$$\begin{aligned}
\tilde{\mu}^+ &:= \sup_{\pi \text{ a p.d. on } \mathcal{D}} \{E(X) \mid \sigma_j^- \leq E_\pi(g_j(X)) \leq \sigma_j^+, j = 1, \ldots, m\}, \\
\tilde{\mu}^- &:= \inf_{\pi \text{ a p.d. on } \mathcal{D}} \{E(X) \mid \sigma_j^- \leq E_\pi(g_j(X)) \leq \sigma_j^+, j = 1, \ldots, m\}, \\
\tilde{\sigma}^2 &:= \sup_{\pi \text{ a p.d. on } \mathcal{D}} \{\text{Var}(X) \mid \sigma_j^- \leq E_\pi(g_j(X)) \leq \sigma_j^+, j = 1, \ldots, m\},
\end{aligned} \tag{9}$$

and assume that $\tilde{\mu}^+, \tilde{\mu}^-, \tilde{\sigma}^2$ are bounded (i.e., this will be the case if $g_j(X) = X^j$, $j = 1, \ldots, m$, with $m \geq 2$ in (1)), and that $\mathcal{D} \subseteq \mathbb{R}^+$ (as in practice). Clearly, for the lognormal distribution mixture (7) to be feasible for the semiparametric bound problem (1), α in (8) should be chosen such that $\alpha^2 \leq \tilde{\sigma}^2$ to ensure that the variance of the lognormal distribution used for the mixture is less than the maximum possible variance of the probability distributions π associated with the expected value constrains in (1). Moreover, as $\alpha \to \tilde{\sigma}$. That is, the only feasible solution of the semiparametric bound problem with lognormal mixture would be a (single) lognormal with variance α^2, and mean x satisfying $\tilde{\mu}^- \leq x \leq \tilde{\mu}^+$, which is unimodal. Thus, there exists an $\alpha \in [0, \tilde{\sigma}]$, such that the lognormal mixture obtained with the CG Algorithm 1 will be unimodal. To find the value of α such that the lognormal mixture obtained with the CG approach is both unimodal and as close as possible to replicate the semiparametric bound obtained by assuming that the probability distribution π in (1) is unimodal (and not necessarily smooth), one can use the *bisection* Algorithm 2.

Algorithm 2 Smooth and unimodal worst/best-case distribution

procedure BISECTION($0 < \alpha_{lo} < \alpha_{hi} < \tilde{\sigma}$, $\epsilon > 0$)
  while $|\alpha_{hi} - \alpha_{lo}| > \epsilon$ **do**
    $\alpha_k \leftarrow \frac{1}{2}(\alpha_{lo} + \alpha_{hi})$
    compute J^*, $p^* := \{p_x^*\}_{x \in J}$, using CG Algorithm 1 and H_x in (7)
    if $\pi \sim \sum_{x \in J^*} p_x^* H_x$ is unimodal **then**
      $\alpha_{hi} \leftarrow \alpha_k$
    else
      $\alpha_{lo} \leftarrow \alpha_k$
    end if
  end while
  return J^*, p^*, $\alpha = \alpha_k$, and M_J^*
end procedure

Note that in the discussion above, the choice of the lognormal distribution is not key. Instead, the same would apply as long as the mixture distribution in (7) is smooth, unimodal, and has at least two appropriate degrees of freedom in the choice of parameters (e.g., in case of random variables

with support on the whole real line, the normal distribution could be used to form the mixture). In Section 5.2, we illustrate with a numerical example how the *bisection* Algorithm 2 can be used to obtain a smooth and unimodal worst-case distribution that closely replicates the behaviour of the worst-case unimodal distribution.

When using a smooth distribution to define the mixture component H_x in (5), it is important to understand the impact of the selection of mixture components H_x. Ideally, computing the bound with a mixture of smooth distributions H_x would yield the optimal value across all possible smooth distributions in the semiparametric bound problem (1). Instead, it is the semiparametric bound across all mixtures with components H_x. However, in Theorem 3 below, we show that the optimal semiparametric bounds across all smooth unimodal distributions is the same as the one across unimodal distributions. Loosely speaking, this follows from the fact that the density function of a uniform distribution can be arbitrarily approximated by an appropriate smooth density function.

Theorem 3. *The semiparametric bound problem (1), with the additional constrain of the underlying distribution being smooth and unimodal is equivalent to problem (1), with the additional constrain of the underlying distribution being unimodal.* (6).

Proof. Let B_u^* be the bound corresponding to the semiparametric bound problem (1), with the additional constrain that the underlying distribution π is unimodal. Note that there exists a distribution π^* such that $B_u^* := E_{\pi^*}(f(X))$ and π^* is a mixture of uniform distributions (cf., Section 4.2); that is, with $H_x \sim \text{Uniform}(\min\{x, M\}, \max\{x, M\})$ in (5). Now, for $\eta > 0$, let π^η be the mixture obtained after replacing H_x by

$$H_x^\eta(u) = \frac{1}{b(x) - a(x)} \left[\frac{1}{1 + e^{-\eta(u-a(x))}} - \frac{1}{1 + e^{-\eta(u-b(x))}} \right], \tag{10}$$

in the mixture π^*, where $a = \min\{x, M\}$ and $b = \max\{x, M\}$, where M is the mode of π^*. The statement follows since by letting $\eta \to \infty$, one obtains a smooth distribution H_x^η (see, Lemma 4 in Appendix A) that is arbitrarily close to H_x (see, Lemma 5 in Appendix A) . □

A numerical example to illustrate Theorem 3 is provided in Section 5.2.

5 Numerical Illustration

Problem (1) is of particular interest in actuarial science because the target function $f(\cdot)$ in (1) can take the form of payoffs for common insurance and risk management products for which the distribution of the underlying random loss is *ambiguous* (see, e.g., Delage and Ye, 2010; Natarajan et al., 2011, for a recent reference); that is, it is not known precisely. Let X represent the loss, and d be the deductible associated with an insurance policy on X. For example, Schepper and Heijnen (2007, Sections 3.1 and 3.2) provide upper and lower bounds on the expected cost per policy payment $\max(X-d,0)$ when only up to third-order moment information on the loss distribution is assumed to be known. This is done by solving (1) analytically with

$$f(X)=\max(X-d,0), g_j(X)=X^j, \sigma_j^+=\sigma_j^- \text{ for } j=1,\ldots,m, \mathcal{D}=\mathbb{R}^+, \text{ and } m=2,3. \tag{11}$$

In practice, losses do not exceed certain maximum, say b. Taking this into account, Cox (1991, Proposition 3.2) provides upper and lower bounds on the expected cost per policy $\max(X-d,0)$ when only the maximum potential loss b and up to second-order moment information on the loss distribution is assumed to be known. Accordingly, this is done by solving (1) analytically with

$$f(X)=\max(X-d,0), g_j(X)=X^j, \sigma_j^+=\sigma_j^- \text{ for } j=1,\ldots,m, \mathcal{D}=[0,b]\subset\mathbb{R}^+, \text{ and } m=2. \tag{12}$$

5.1 Second-order LER Bounds with Unimodality

Let us reconsider the semiparametric bound on the expected cost per policy defined in (12). Note that from semiparametric bounds on the expected cost per policy, one can readily obtain bounds on the expected *loss elimination ratio* (LER) of the policy. Specifically, note that the expected LER associated with a policy with payoff $\max\{0,X-d\}$ is (cf., Cox, 1991)

$$E(\mathrm{LER}(X))=\frac{E(\min(X,d))}{E(X)}=\frac{E(X)-E(\max(X-d,0))}{E(X)}. \tag{13}$$

Being able to compute bounds on the expected LER would be beneficial for an insurer attempting to set a deductible in cases where the actual loss distribution is ambiguous. For example, in Cox (1991, Section 3), the relationship (13) is used to obtain upper and lower semiparametric bounds on

the expected LER of an insurance policy with deductible, assuming only the knowledge of the mean and variance of the loss, and that the loss cannot exceed a known maximum.

Another sensible premise is to assume that the loss distribution is unimodal. To illustrate the potential of the CG approach, in Figure 1, we use the mixture transformation of Section 4.2 to compute upper and lower semiparametric bounds on the insurance policy with payoff $\max\{0, X - d\}$ when the mean μ and variance σ^2 of the loss are assumed to be known, and the loss cannot exceed the value of b, where d is the policy's deductible, and the loss distribution is also assumed to be unimodal with mode M . The results are compared with the analytical formula of Cox (1991, Section3) to illustrate the tightening of the bounds obtained by adding the unimodality assumption. Specifically, following Cox (1991, Section 4), in Figure 1 we set $\mu = 50$, $\sigma = 15$, $b = 100$, and $M = \{45, 50\}$.

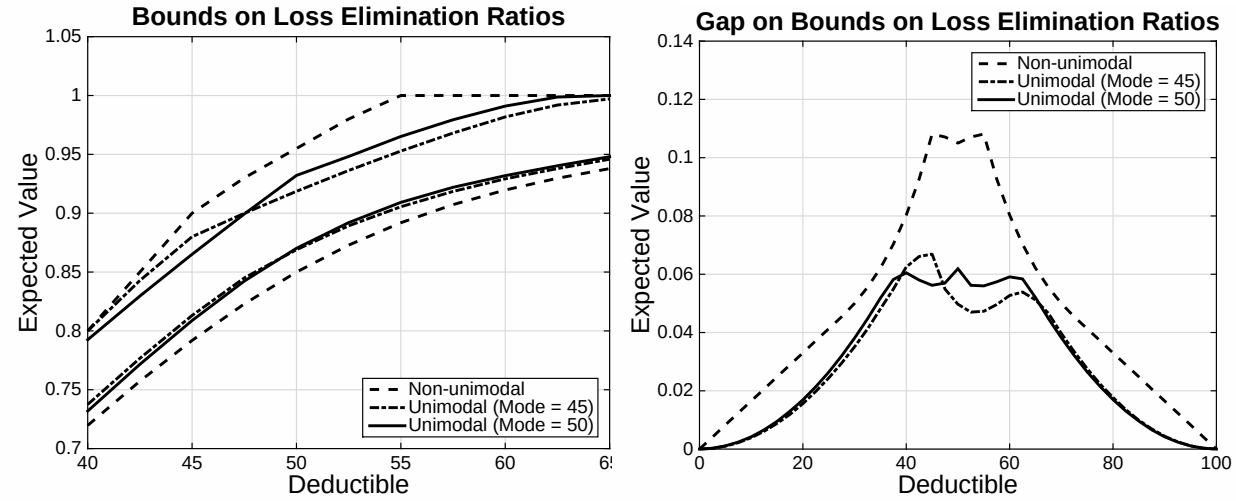

Figure 1: Expected LER bounds (left) and gaps (right) for different values of the deductible d, when the mean $\mu = 50$, and variance $\sigma^2 = 225$ of the underlying loss, as well as its potential maximum value $b = 100$, are assumed to be known. Gaps indicates the difference between upper and lower bounds. Results are presented for bounds without the unimodality constrain, and with unimodality constrain with mode $M = \{45, 50\}$.

Observe in Figure 1 that the expected LER *gap*; that is, the difference between the upper and lower semiparametric expected LER's bounds, is significantly tighter when unimodality is included. When $M = 50$ the gap is symmetrical with a small spike at the mean/mode. The case in which $M = 45$ yields a corresponding peak in the gap around the mode. For either very high or very low deductible values, the choice of the mode has little impact on the size of the gap. Regardless of the mode's value, the gaps are of similar magnitude and narrower than in the case where the underlying loss distribution is not assumed to be unimodal.

5.2 Examining worst-case distributions

Suppose we wish to compute the semiparametric upper bound defined by (11) with $m = 2$. Specifically, let the expected value (moment) constrains in (1) be given by

$$E_\pi(X) = \mu, E_\pi(X^2) = \sigma^2 + \mu^2, \tag{14}$$

for given $\mu, \sigma \in \mathbb{R}^+$. That is, both the mean and the variance of the underlying loss distribution are assumed to be known. In this case, equation (9) reduces to $\mu_\pi^+ = \mu_\pi^- = \mu$, and $\sigma_\pi^2 = \sigma^2$.

A closed-form solution for the corresponding semiparametric upper bound problem was derived by Lo (1987), where he considers f in (11) as the payoff of an European call option with strike d. If furthermore, it is assumed that the underlying distribution π of the loss (or asset price) is unimodal, the corresponding semiparametric upper bound can be computed using: the analytical formula provided by Schepper and Heijnen (2007, Section 3.3), the SDP techniques provided by Popescu (2005), or the uniform mixture approach of Section 4.2 with components (6). The CG uniform mixture method readily provides a worst-case distribution. This distribution, however, is not smooth, has finite support, and is unrealistic as a model for the uncertainty of losses. For this reason we compute upper bounds using smooth mixture compontents in (5) and inspect the resulting worst-case probability and cumulative density functions. The resulting smooth distribution is then compared to the optimal unimodal uniform mixture distribution. Specifically, we use the lognormal mixture (7).

In particular, let us sample the values of μ, σ in (14) from a lognormal asset price dynamics model which is also commonly used to model (a non-ambiguous) loss distribution (see, e.g. Cox et al., 2004; Jaimungal and Wang, 2006). Namely, let $\mu = X_0 e^{rT}$, and $\sigma = X_0 e^{rT}(e^{\nu^2 T} - 1)^{\frac{1}{2}}$ for values of $X_0 = 49.50$, $r = 1\%$, $\nu = 20\%$, and $T = 1$. Also let $d = X_0$ in (11); that is, consider a policy where the expected value of the loss is equal to the deductible.

The semiparametric upper bound was computed using the lognormal mixture (7) for different values of $\alpha \in [1, 1.5, \ldots, 20]$ and the percent above the parametric value of the policy based on Black-Scholes formula was plotted in Figure 2. The corresponding semiparametric bound without the unimodality assumption (given by Lo (1987)), and unimodal bounds with an uniform mixture from Section 4.2 are also plotted for reference. The bold point in Figure 2 represents the smooth,

unimodal bound obtained with the bisection Algorithm 2 and a mixture of lognormal distributions. Also, the plot illustrates the point ($\alpha = 11.34$) at which the the bounds obtained by the bisection Algorithm 2 and the CG Algorithm 1 with a mixture of uniform distributions are equal.

In Figure 2 one can observe that for extremely low values of α, the component distributions of the mixture are very narrow, approaching the pessimistic discrete distribution case associated with closed-form bound of Lo (1987). We also see that as $\alpha \to \sigma = 20.4$ the resulting bound distribution converges to the lognormal specified by the Black and Scholes asset pricing framework. This convergence is seen in Figure 2 since the error goes to zero and the upper bound price equals the analytical Black-Scholes price.

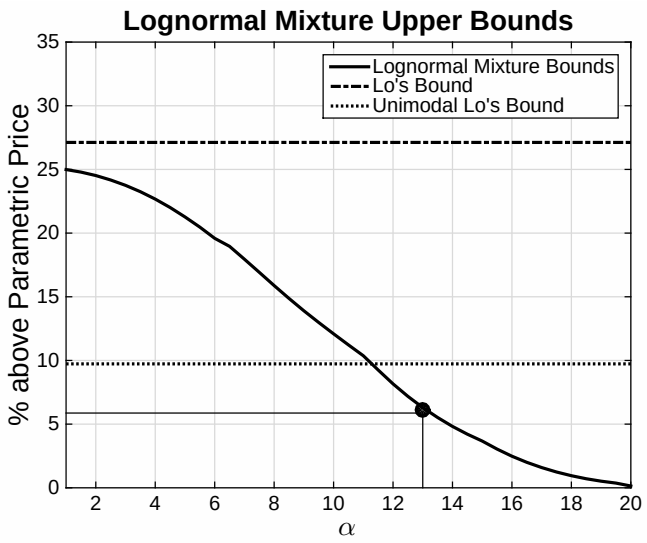

Figure 2: Percentage above the parametric Black and Scholes price of the Lo (1987) upper bound (Lo's Bound) and the lognormal mixtures obtained from Algorithm 2 (lognormal Mixture Bounds). The bold point denotes the value of $\alpha = 13.75$ in which the lognormal mixture bound obtained from Algorithm 2 produces a unimodal distribution.

Figure 2 also highlights the result discussed in Theorem 3. The bound computed using uniform mixture components is greater than the unimodal bound from the lognormal mixture with the gap size under 4%. Note that the unimodal upper bound using lognormal mixture components occurred at $\alpha = 13.75$. As mentioned before, the α that yields the same bound value as that from the uniform mixture is $\alpha = 11.34$. The smaller the α in the lognormal mixture (7) the higher the conservatism associated with the semiparametric bound. Figure 4 shows the probability distribution function (PDF) and cumulative distribution function (CDF) of the lognormal mixtures at $\alpha = \{11.34, 13.75\}$ as well as the associated *true* lognormal distribution with mean and variance equal to the moments

used to compute the semiparametric bounds. To highlight the advantage of using the lognormal mixtures instead of the uniform mixtures, Figure 3 shows the optimal PDF and CDF of the latter along with the associated true lognormal distribution.

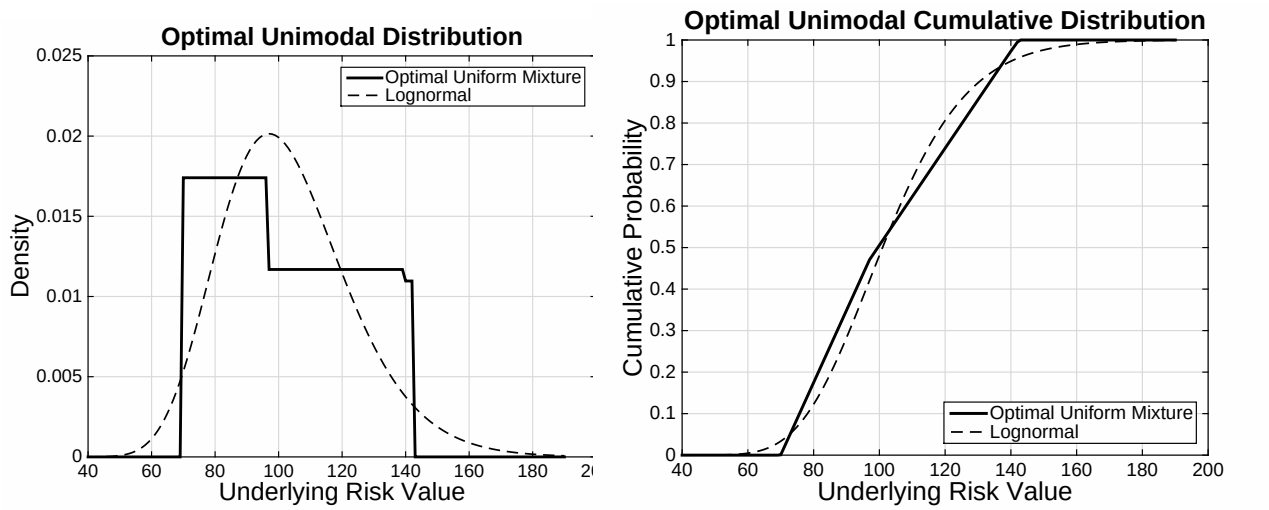

Figure 3: PDF and CDF that yields the optimal unimodal bound via uniform mixtures compared with an associated lognormal distribution.

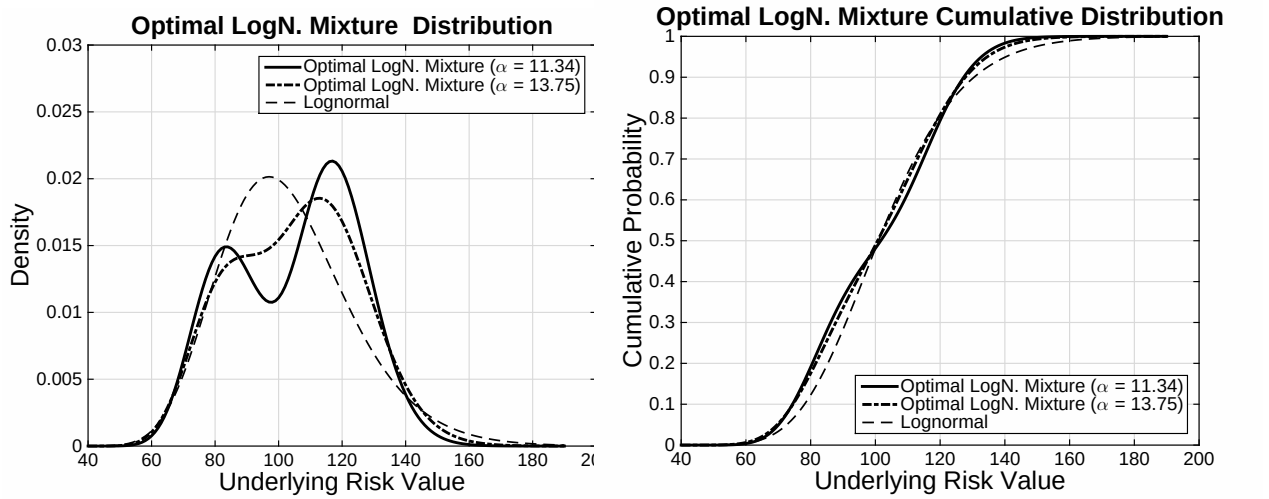

Figure 4: PDFs and CDFs that yield the optimal bounds via lognormal mixtures (cf., Algorithm 2) for $\alpha = \{11.34, 13.75\}$, compared with an associated lognormal distribution.

Observe in Figure 4 that the unimodal lognormal mixture at $\alpha = 13.75$ is relatively close to the shape of the associated true lognormal distribution. Contrast this with the PDF of the unimodal mixture of Figure 3 which bares little similarity to the associated true lognormal probability density. The primary difference to note is that the lognormal mixture approach yields an unimodal distribution, but the mode is not specified. Using the uniform mixture approach of Section 4.2 will produce a

density with a specified mode, but at the cost of an unrealistic distribution. The lognormal mixture at $\sigma = 11.34$ is bimodal and does not resemble the true density. In each case the cumulative densities are fairly close to the true CDF. This example highlights the ability of the lognormal mixture approach to construct realistic unimodal distributions while still being close to the optimal unimodal bound; here the gap was shown to be under 4%.

5.3 Illustration of Theorem 3

We finish this section by providing numerical results to illustrate Theorem 3. Reconsider the semiparametric bound problem defined in (11) with $m = 2$ (i.e., with up to second-order moment information), and the additional constrain that the underlying loss distribution is unimodal.

Suppose we compute the semiparametric upper bound of the at-the-money policy described in Section 5.2 enforcing the first two known moments and continuity. To illustrate Theorem 3, the upper bound is also computed for the option using (10) and various levels of η. The percentage difference between the former and latter are plotted against η in Figure 5.

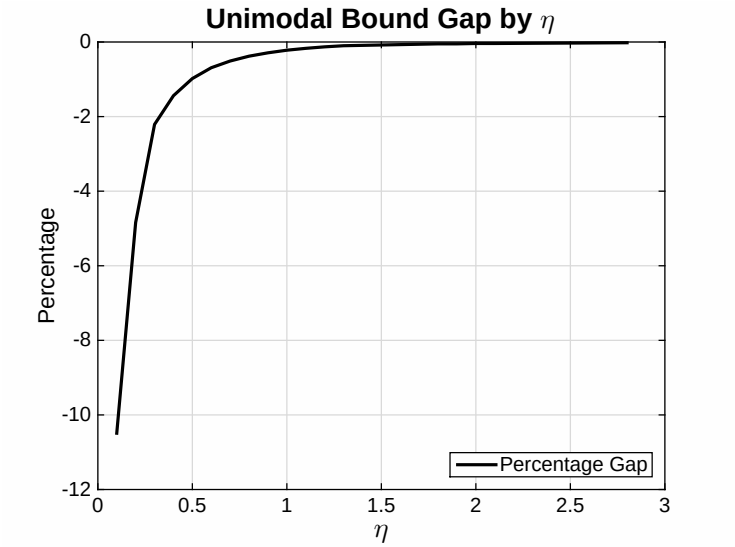

Figure 5: Illustration of how the upper bound with mixture components (10) converges to the unimodality bound (6) (without smoothness requirements) as $\eta \to \infty$. The plot shows the difference in percentage between these bounds as a function of η.

From Figure 5 we see that by implementing the algorithm with (10) and increasing η the upper bound converges to that computed with mixture component (6). Bounds computed using smoothness and unimodality can yield values arbitrarily close to, but not greater than those obtained when only unimodality is enforced. The implication of Theorem 3 is that any tightening of the upper bound

from a smooth mixture is a byproduct of the choice of the mixture distribution H_x, not from the inclusion of smoothness. In practice it can be confirmed that the change in bound from choice of H_x is generally fairly small.

6 Extensions

Besides the common features of insurance policies considered in Section 5 such as the policy's deductible d, and the fact that losses typically do not exceed a maximum loss b; a maximum payment, and coinsurance are also common features in insurance policies. If a policy will only cover up to a maximum loss of $u \in R^+$, and coinsurance stipulates that only some proportion $\gamma \in [0, 1]$ of the losses will be covered, then the policy's payoff can be written as $f(X) = \gamma[\min(X, u) - \min(X, d)]$. All of these policy modifications can be readily incorporated into the CG solution approach framework. In particular, the CG methodology can be applied to compute bounds on the expected policy loss, for a wide variety of standard functions of loss random variables used in industry.

In the numerical examples in Section 5, the target function f in (1) was used to model piecewise linear insurance policy payoffs, and the functions g_j, $j = 1, \ldots, m$ to use the knowledge of up to m-order non-central moments of the loss distribution. However, the methodology discussed here applies similarly to functional forms of f that are not piecewise linear policy payoffs, and the functions g_j, $j = 1, \ldots, m$ can be used to represent the knowledge of other than non-central moment's information. As an example, let c_j be the European call prices on some stock X for various strike prices K_j, $j = 1, \ldots, m$. Recall that the payoff for this type of option is the same as that of the d-deductible policy described in (11). The constrain set for (1) can be defined to enforce market prices of options by setting

$$E_\pi(g_j(X)) := E_\pi(\max(X - K_j, 0)) = c_j, j = 1, \ldots, m. \tag{15}$$

The market price constrains (15) can then be used to compute semiparametric bounds on the variance of the underlying asset. This is accomplished by definining the target function $f(X)$ in (1) as

$$E_\pi(f(X)) := E_\pi((X - \mu)^2) \tag{16}$$

where μ is the known first moment of X. In Bertsimas and Popescu (2002) it was shown that com-

puting semiparametric bounds on (16) using knowledge on the prices (15) is a useful alternative to the standard methods of computing implied volatilities. For risk management purposes, semiparametric bounds can also be used to compute bounds on one sided deviations of the underlying risk, that is, its semivariance. Each of these common applications readily fit into the framework of the CG methodology presented in Section 3, demonstrating the variety of contexts in which the CG approach can be used to compute semiparametric bounds.

7 Summary

The CG methodology presented here provides a practical optimization-based algorithm for computing semiparametric bounds on the expected payments of general insurance instruments and financial assets. Section 3 described how the general problem described in (1) can be solved, in most practical instances, by solving a sequence of linear programs that are updated using simple arithmetic operations. The CG approach also readily provides a representation of the worst/best-case distributions associated with a semiparametric bound problem.

To illustrate the potential of the of the CG algorithm semiparametric bounds on the payoff of common insurance policies were computed. It was also shown that additional distributional information such as continuity and unimodality can be incorporated in the formulation in a straightfoward fashion. The ability to include these constrains provides tighter bounds on the quantity of interest as well as distributions that fit the practitioner's problem specific knowledge. Note that from the recent work of Lee and Lin (2010), it follows that for some Property/Casualty insurance problems it will be suitable to consider that the underlying random variable follows a distribution that is a mixture of Erlang distributions. The advantage of using mixtures of Erlang distributions is the existence of extremely efficient expectation–maximization (EM) algorithms for parameter estimation from raw data. This interesting line of work will be the subject of future research work.

The CG methodology offers a powerful and compelling alternative for computing semiparametric bounds in comparison with the main approaches used in the literature to compute them; namely, deriving analytical solutions to special cases of the problem or solving it numerically using semidefinite programming. This is due to the speed, generality, and ease of implementation of the CG algorithm. The CG algorithm achieves accurate results at a very small computational cost. The straightforward

implementation used for the test problems shown here generated solutions in at most 1-2 seconds. Furthermore, although the examples considered here focused on obtaining semiparametric bounds for insurance policies with piecewise linear payoff given moment information about the underlying loss, the CG approach presented here allows for a very general class of univariate semiparametric bounds to be computed using the same core the solution algorithm.

Acknowledgements

This work has been carried out thanks to funding from the Casualty Actuarial Society, Actuarial Foundation of Canada, and Society of Actuaries.

Appendices

A Smooth Approximations of the Uniform Distribution

Let X be a random variable that is uniformly distributed on the interval $[a, b]$. The probability density function (PDF) of X is $f(x) = \frac{1}{b-a}$. It is possible to construct a smooth function that approximates $f(x)$ and is asymptotically equal to the true PDF. We define the following η parameterized function.

$$f_\eta(x) = \frac{1}{b-a}\left[\frac{1}{1+e^{-\eta(x-a)}} - \frac{1}{1+e^{-\eta(x-b)}}\right] \tag{A.1}$$

The probability function $f_\eta(x)$ is the difference in two shifted logistic functions.

Lemma 4. *For any $\eta > 0$, and $a, b \in \mathbb{R}$ such that $b \geq a$, the cumulative probability distribution $F_\eta(x)$ associated with the probability distribution $f_\eta(x)$ is $F_\eta(x) = 1 - \frac{1}{\eta(b-a)} \ln\left(\frac{1+e^{-\eta(x-b)}}{1+e^{-\eta(x-a)}}\right)$. In particular, $\lim_{x\to\infty} F_\eta(x) = 1$, $\lim_{x\to-\infty} F_\eta(x) = 0$, and $\frac{dF_\eta(x)}{dx} \geq 0$ for all $x \in \mathbb{R}$.*

In Lemma 5 we show that as $\eta \to \infty$ (A.1) will converge to the PDF of X.

Lemma 5. *As $\eta \to \infty$ the function $f_\eta(x)$ in* (A.1) *converges to a uniform distribution on $[a, b]$.*

Proof. To show that $f_\eta(x) \to \text{Uniform}(a, b)$ as $\eta \to \infty$ consider three different cases corresponding to three intervals of x. First consider the case in which $x < a$. For $x < a$ each of the exponent terms are positive, i.e. $0 < -\eta(x-b) < -\eta(x-a)$, for all $\eta > 0$. So, for $x < a$ we see that $f_\eta(x) \to 0$. Next we

look at $x > b$. In this case each of the exponent terms are negative, which again yields a limit of 0. Finally, consider $a < x < b$. On this interval the exponent terms satisfy $-\eta(x-a) < 0 < -\eta(x-b)$ for all $\eta > 0$. So, for $a < x < b$ we have $f_\eta(x) \to \frac{1}{b-a}(1-0) = \frac{1}{b-a}$. We can conclude that the limit of $f_\eta(x)$ is as follows

$$\lim_{\eta\to\infty} f_\eta(x) = \begin{cases} 0 & x < a \\ \frac{1}{b-a} & a < x < b \\ 0 & x > b \end{cases}$$
$$= \text{Uniform}(a, b)$$

□

To demonstrate Lemma 5, consider $X \sim \text{Uniform}(20, 30)$. The PDF of X can be approximated using (A.1) with $a = 20$ and $b = 30$. In Figure 6 we plot the PDF of X as well as the approximation curve for different values of η.

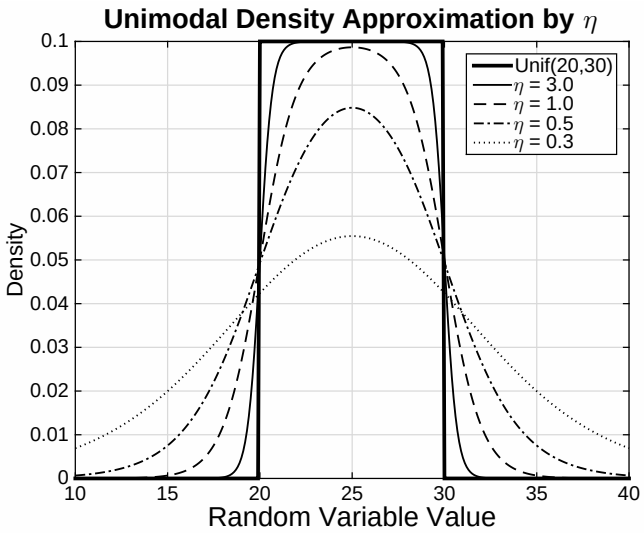

Figure 6: PDF of Uniform(20, 30) distribution along with the approximating density (A.1) for different values of η.

From Figure 6, we see that as the parameter η increases the curve of (A.1) approaches the PDF of X.

References

Bertsimas, D. and Popescu, I. (2002). On the relation between option and stock prices. *Operations Research*, 50(2):358–74.

Bertsimas, D. and Tsitsiklis, J. N. (1997). *Introduction to Linear Optimization*. Athena Scientific, Belmont, Massachusetts.

Birge, J. R. and Dulá, J. H. (1991). Bounding separable recourse functions with limited distribution information. *Annals of Operations Research*, 30:277–298.

Boyle, P. and Lin, X. (1997). Bounds on contingent claims based on several assets. *J. Fin. Econ.*, 46(3):383–400.

Brockett, P., Cox, S., and Smith, J. (1996). Bounds on the price of catastrophe insurance options on futures contracts. *Securitization of Insurance Risks. Society of Actuaries Monograph Series*.

Chen, L., He, S., and Zhang, S. (2011). Tight bounds for some risk measures, with applications to robust portfolio selection. *Operations Research*, 59(4):847–865.

Cox, S. (1991). Bounds on expected values of insurance payments and option prices. *Transactions of the Society of Actuaries*, XLIII:231–60.

Cox, S., Lin, Y., Tian, R., and Zuluaga, L. (2010). Bounds for probabilities of extreme events defined by two random variables. *Variance*, 4(1):47–65.

Cox, S., Lin, Y., Tian, R., and Zuluaga, L. F. (2011). Mortality portfolio risk management. *Journal of Risk and Insurance*.

Cox, S. H., Fairchild, J. R., and Pedersen, H. W. (2004). Valuation of structured risk management products. *Insurance: Mathematics and Economics*, 34:259–272.

Dantzig, G. B. (1963). *Linear Programming and Extensions*. Princeton University Press. `http://www.rand.org/pubs/reports/R366.html`.

Dantzig, G. B. and Wolfe, P. (1961). The decomposition algorithm for linear programming. *Econometrica*, 29(4).

Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612.

Jaimungal, S. and Wang, T. (2006). Catastrophe options with stochastic interest rates and compound Poisson losses. *Insurance: Mathematics and Economics*, 38:469–483.

Jansen, K., Haezendonc, J., and Goovaerts, M. (1986). Upper bounds on stop-loss premiums in case of known moments up to the fourth order. *Insurance: Mathematics and Economics*, 5:315–34.

Kemperman, J. H. B. (1968). The general moment problem, a geometric approach. *The Annals of Mathematical Statistics*, 39(1):93–122.

Khintchine, A. Y. (1938). On unimodal distributions. *Izv. Naucho-Isled. Inst. Mat. Mech. Tomsk. Gos. Univ.*, 2:1–7.

Lee, S. and Lin, S. (2010). Modeling and evaluating insurance losses via mixtures of erlang distributions. *North American Actuarial Journal*, 14(1).

Lo, A. (1987). Semi-parametric upper bounds for option prices and expected payoffs. *Journal of Financial Economics*, 19(2):373–388.

Lubbecke, M. E. and Desrosiers, J. (2005). Selected topics in column generation. *Operations Research*, 53(6):1007–023.

Natarajan, K., Sim, M., and Chung-Piaw, T. (2011). Beyond risk: Ambiguity in supply chains. In *Handbook of Integrated Risk Management in Global Supply Chains*, pages 103–124. John Wiley & Sons Inc.

Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer, New York, 2nd edition.

Popescu, I. (2005). A semidefinite programming approach to optimal moment bounds for convex classes of distributions. *Mathematics of Operations Research*, 30(3):632–57.

Schepper, D. and Heijnen, B. (2007). Distribution-free option pricing. *Insurance: Mathematics and Economics*, 40:179–99.

Tawarmalani, M. and Sahinidis, N. V. (2005). A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103:225–249.

Todd, M. (2001). Semidefinite optimization. *Acta Numerica*, 10.

Villegas, M., Medaglia, A. L., and Zuluaga, L. (2010). Computing bounds on the expected payoff of alternative risk transfer products. *Mathematics: Insurance and Economics (forthcoming)*.

Zuluaga, L. and Peña, J. (2005). A conic programming approach to generalized tchebycheff inequalities. *Mathematics of Operations Research*, 30(2):369–88.